

Quarterly Progress Report #8

**Prepared by John Donoghue, Nicholas Hatsopoulos, Mijail Serruya,
Elie Bienenstock, Lauren Lennox and Beth Travers**

Quarter 8: July 1, 2001-September 30, 2001

Neural Prosthetic Control

NINDS CONTRACT NUMBER: N01NS92322

Department of Neuroscience and Brain Science Program

Brown Medical School

Brown University

Contents

1. Introduction	2
2. Summary of Related Achievements this Quarter	3
3. Neural Decoding	3
3.1 Non Parametric decoding algorithms	4
3.2 Principal Component Analysis	5
3.3 Tuning Maps	6
3.4 Wilcoxon Signed Rank Test	6

This Quarterly Report will focus on progress made in Neural Decoding. In what follows, we report on the estimation of conditional firing maps. Each such map (for a given cortical cell) can be thought of as a 2-D *tuning function*, or *receptive field*, which characterizes the *response* of the cell given hand velocity or hand position. It is one of the two major components of the generative model that we use in our Bayesian approach, and needs to be estimated (learned) off-line. First there is a summary of the overall objectives of the contract.

1. Introduction

A number of neurological disorders, such as spinal cord injury, MD and ALS result in the inability to make voluntary movements. A major reason for paralysis in these disorders is a disconnection of the signal from a normal brain from the spinal cord or muscles. Devices that can detect and decode motor commands have the potential to restore voluntary actions in these individuals. The purpose of this project is to demonstrate the ability to use neural signals to control real world devices in monkeys; such devices can ultimately serve as prosthetic aids for paralyzed individuals.

Control signals for prosthetic devices can be derived from a number of sources, including the eyes, muscles, and EEG. These signals are, however, rather limited in the number of dimensions they can control. Going beyond a one dimensional control signal is difficult and often interferes with natural behavior. For example, two dimensional EEG control requires full attention to control without distraction (such as gaze shifts). By contrast, populations of neurons appear to contain rich signals, potentially able to control multiple dimensions independently. However, chronic recording of multiple neurons in primates has been technically challenging, the ability to decode neural activity into meaningful control signals is poorly understood and the ability to control devices using such signals is not fully developed.

The overall goal of this work is to develop a means to bring a robotic arm under near real time neural control using a multineuron signal derived from a recording device that is chronically implanted in a macaque monkey motor cortex. This project has three specific objectives. The **first objective** is to develop and test technologically advanced neural recording devices in a non-human primate model. This work examines the stability, efficiency and biocompatibility of electrode arrays and the suitability of the primary motor

cortex as a sight to obtain neural recordings. Once recorded, neural activity must be decoded into meaningful control signals. The optimal methods for such decoding are not obvious. A **second objective** of the project is to examine various decoding methods and evaluate their ability to be useful control signals. This requires mathematical tools and signal processing that reconstructs intended actions from abstract, neurally based motor commands generated in the cortex. This aspect of the project involves fundamental motor control questions, such as what coordinate system is used to encode voluntary actions. A **third objective** of this project is to show that such signals can be used to control devices such as a robotic arm or a computer interface. These devices serve as a proxy for the lost limb and can be used to recreate useful actions like those intended for the arm. Successful completion of these goals would suggest that this approach could be used to restore movement in paralyzed humans.

2. Summary of Related Achievements this quarter

This quarter we implanted one additional array (99-3) in MI and further developed methods to enhance the speed of training monkeys for use in array testing. We sent out 4 hemispheres for histology (Thionin and GFAP staining). We continued development of neural decoding methods, including new probabilistic methods.

3. NEURAL DECODING

The goals of this aspect of the project are to determine: whether we can recover the hand trajectory using the activity of multiple MI neurons; how reliable this reconstruction will be; how many simultaneously recorded neurons are required; can the computation be performed fast enough to be used in a prosthetic device; and finally can the reconstruction algorithm be made adaptive enough that it will withstand changes in the functional properties of recorded neurons as may result for instance from motion of the implant between successive days or weeks (see above). In a first step, we have demonstrated that linear regression methods based upon small numbers of neurons provide a moderately accurate estimate of any new hand trajectory. This work is being submitted for publication. In a second step we developed non-parametric Bayesian methods with the goal of achieving better trajectory reconstructions. The provisional conclusions of our study at the present time, based on partly simulated data, suggest that simultaneous from several hundred cells, when they are available, will provide very accurate reconstructions of hand trajectory.

In previous works, authors have considered a variety of models, including a cosine tuning function (Georgopoulos et al. 1986), and a modified cosine function (Moran and Schwartz 1999). In these highly-constrained, *parametric*, models, smoothness is intrinsically built in, in the form of strong assumptions about the shape of the tuning function.

3.1 Non-Parametric Decoding Algorithms In contrast, we have explored various *non-parametric* models, which can accommodate a wide range of functional forms, and where smoothness can be adjusted. Hand position (x and y position of manipulandum, digitized at 167 Hz) is first smoothed using a smoothing spline, and then subsampled at regular intervals of 50 ms. We then compute the x and y velocities as the derivatives of the spline at these knot points, and convert them to polar coordinates (r, θ) , where r is the speed and θ is the direction of motion. Spike trains are binned using the same 50-ms resolution, and

in the following, we report on maps for the spike count in the $[t-150, t-100]$ bin, where t (ms) is the time at which the velocity is measured. The velocity space (polar coordinates) is discretized, using a 100x100 grid. We denote by $f(v)$ the empirical mean firing of a given cell at velocity v . This empirical estimate is built from a training set representing about 15 minutes worth of data (training data is limited to recording periods that satisfy several requirements meant to ensure stationarity to the extent possible). The sampling of the 100x100 velocity grid is nonhomogeneous and relatively sparse in some regions. Clearly, since the data is noisy and sparse, we need to compute an optimal estimate of the tuning function by an appropriate smoothing of the training data. We shall denote this estimate by $g(v)$.

Our non-parametric models are related to Markov Random Fields (MRF) (Geman and Geman 1984), and include a spatial *prior probability*, which encodes our expectations about the variation of neural activity in velocity space. The MRF prior states that the expected firing at a given velocity depends only on the firing at the 4 neighboring velocities in the discretized 100x100 grid. We consider two possible prior models: Gaussian and "robust." A Gaussian prior corresponds to the assumption that the firing rate varies smoothly. A robust prior (e.g. a Student's t-distribution) assumes a heavy-tailed distribution of the spatial variation, and implies piecewise smooth data. The use of such priors is motivated by the examination of the histograms of differences of observed firing rates between adjacent velocities. These histograms (or log-histograms) are heavy-tailed and exhibit shapes typical of images of natural scenes, where robust statistical error functions are often used.

Our models also include a term that represents the *likelihood* of observing a particular firing rate $f(v)$ given the true (i.e., to-be-estimated) rate $g(v)$. This likelihood corresponds to a particular *generative* model for the spike count $f(v)$ in a 50-ms time bin, given the mean $g(v)$. A convenient—albeit biologically somewhat implausible—form for the generative model is a Gaussian distribution. A biologically more satisfactory model is the Poisson distribution of mean $g(v)$.

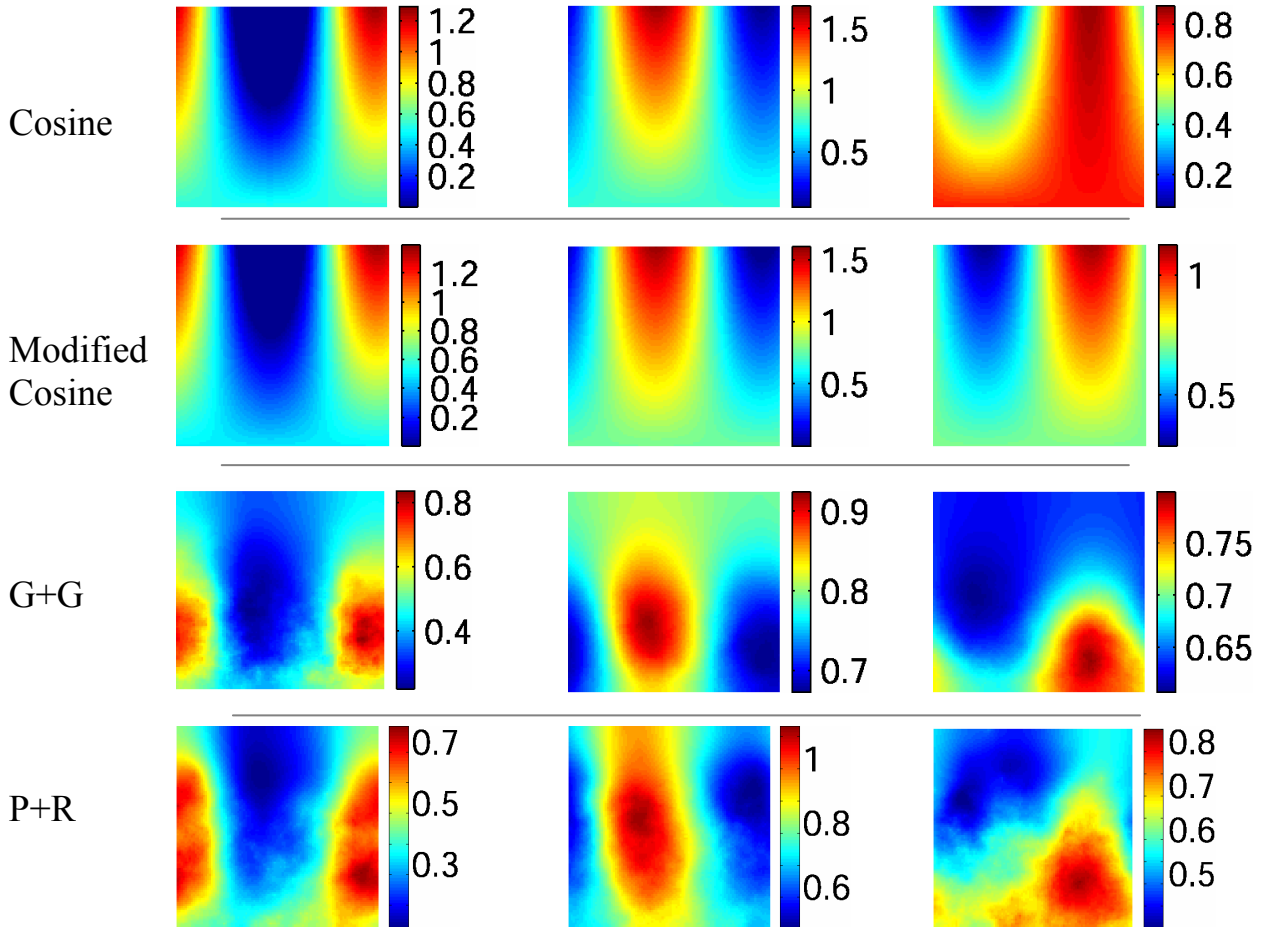
Adopting a Bayesian formulation, we construct *Maximum A Posteriori* (MAP) estimates of a cell's conditional firing $\mathbf{g} = \{g(v)\}$. In the special case of Gaussian prior and Gaussian likelihood (we refer to this case as the G+G model below), the MAP estimate is easily seen to be the minimizer of a quadratic *energy function*, with a parameter λ which adjusts the amount of smoothing. When λ is 0, the optimal \mathbf{g} is the observed firing rate at each v . If on the other hand λ is infinite, the optimal \mathbf{g} is a constant function, equal to the global mean of the observed firing rates. The smoothing parameter λ is adjusted by cross-validation. This is a standard regularization problem, which can be solved in closed form. In contrast, when using a Poisson likelihood (generative model) and Robust prior (referred to as P+R below), no closed-form solution exists. We then find a "reasonable" (possibly sub-optimal) solution by an iterative algorithm.

3.2 Principal Component Analysis We further examined the result of applying Principal Component Analysis (PCA) to our non-parametric estimates of the tuning functions for a collection of 25 motor cortical cells. Each tuning function is a 10,000-

dimensional vectors, and so are the eigenvectors. We typically use the first 4 largest eigenvalues, which account for about 80% of the variance.

To obtain a quantitative comparison of various models, we used a cross-validated log-likelihood criterion. Specifically, 10 trials out of 180 were left out for testing, and models were fit on the remaining data, yielding estimated conditional mean rates ("firing maps"). We then computed the log-likelihood of the spike counts in the test data, given the model. This provides a measure of how well the model captures the statistical variation in the training set. The whole procedure was repeated 18 times for different test/training partitions.

3.3 Tuning Maps The following figure shows examples of tuning maps for three different cells under the four different models: Cosine (Georgopoulos et al. 1986), Modified Cosine (Moran and Schwartz 1999), Non-Parametric G+G, Non-Parametric P+R.



3.4 Wilcoxon Signed Rank Test To test whether one model is significantly better than another, we applied a nonparametric test: the Wilcoxon signed rank test. The following table shows that the non-parametric models do a better job of explaining new data than the parametric models. The Poisson+Robust fit provides the best description of the data, and PCA post-processing yields further significant improvement. Note the very high p -values. The improvement afforded by the P+R model indicates that the conditional firing rate is well described by regions of smooth activity with relatively sharp discontinuities between them. It appears that PCA reduces the variance of nonparametric models without increasing much of the bias, and so it increases the log-likelihood significantly.

Method	Log Likelihood Ratio	p-value
G+G over Cosine	24.9181	7.6294e-06
G+G over M/S	15.8333	0.0047
P+R over Cosine	50.0685	7.6294e-06
P+R over M/S	32.2218	7.6294e-06
PCA of P+R over P+R	233.6086	7.6294e-06